

AMENDMENTS TO THE CLAIMS

1. (Currently Amended) A distributed file storage system comprising:
 - a plurality of storage units configured to communicate with each other;
 - said plurality of storage units including:
 - a first storage unit including a storage disk and a processor;
 - a second storage unit including a storage disk and a processor;
 - a third storage unit including a storage disk and a processor; and
 - a fourth storage unit including a storage disk and a processor;
 - a file stored on the distributed file storage system;
 - a first file portion of the file comprising a first set of file data stored in the first storage unit;
 - a second file portion of the file comprising a second set of file data stored in the second storage unit, wherein the second set of file data is different from the first set of file data;
 - a first metadata to identify in part the location of the file, the first metadata stored on the first storage unit, the second storage unit, the third storage unit, and the fourth storage unit;
 - a second metadata, different at least in part from the first metadata, to supplement the first metadata in identifying the location of the file, the second metadata stored on at least one, but not all, of the first storage unit, the second storage unit, the third storage unit, and the fourth storage unit;
 - a switch in communication with the plurality of storage units, the switch configured to receive a read request for the file stored on the distributed file storage system and to send the read request to one of the plurality of storage units wherein each of the plurality of storage units is operable to monitor in real time a pattern of access to the file, a latency to access each copy of the file, and content included in the file, such that a block cache module will perform K packet read aheads, where K is calculated using a current read rate and a current latency of an access link; and
 - each of the plurality of storage units is configured to use the first metadata to process a read request on behalf of the distributed file storage system, wherein the distributed file storage system is arranged for dynamically determining at least one copy of the file to be replicated and dynamically determining a quantity of the plurality of storage units to store each replicated copy of the file based

at least in part on the real time monitoring of the pattern of access to the file, the latency to access each copy of the file, and content included in the file.

2. (Previously Presented) The distributed file storage system of Claim 1, further comprising error correction data related to the file data, the error correction data stored in the distributed file storage system.

3. (Previously Presented). The distributed file storage system of Claim 2, further comprising error correction data which includes parity information.

4. (Previously Presented) The distributed file storage system of Claim 3, further comprising parity information which includes parity data blocks and location information indicating where the parity data blocks are stored.

5. (Previously Presented) The distributed file storage system of Claim 2, further comprising error correction data which includes redundancy data, wherein the second metadata indicates the location of the redundancy data.

6. (Previously Presented) The distributed file storage system of Claim 1, further comprising first metadata which includes metadata related to the root directory.

7. (Previously Presented) The distributed file storage system of Claim 1, the plurality of storage units comprising storage units configured to receive a request and to initiate the request to move the first file portion in real-time from the first storage unit to the third storage unit, and to send a request to update the second metadata to indicate the location of the moved first file portion.

8. (Previously Presented) The distributed file storage system of Claim 1, the plurality of storage units comprising storage units configured to receive a request and to initiate the request to replicate the first file portion in real-time and to store the replicated first file portion on a different storage unit, and to send a request to update the second metadata to indicate the location of the replicated first file portion.

9. (Previously Presented) The distributed file storage system of Claim 1, further comprising second metadata which includes metadata related to the locations in which the file data is stored.
10. (Previously Presented) The distributed file storage system of Claim 1, further comprising second metadata which includes metadata related to a parent directory of the file.
11. (Previously Presented) The distributed file storage system of Claim 1, further configured to handle more READ requests than WRITE requests.
12. (Previously Presented) The distributed file storage system of Claim 1, further configured to handle block transactions.
- 13-42. (Canceled)
43. (Previously Presented) The distributed file storage system of Claim 1, wherein the file has been stored on a number of the plurality of storage units, wherein the number is determined specifically for the file, and wherein the number is equal to or greater than two.
- 44-70. (Canceled)
71. (Previously Presented) The distributed file storage system of Claim 1, wherein the first metadata identifies the location of the second metadata.
72. (Previously Presented) The distributed file storage system of Claim 1, wherein the second metadata identifies the location of a third metadata, different at least in part from the first metadata and the second metadata, to supplement the first metadata and the second metadata in identifying the location of the file.
73. (Previously Presented) The distributed file storage system of Claim 1, wherein the second metadata identifies the location of at least one of the following: the first set of file data and the second set of file data.
74. (Canceled)

75. (Currently Amended) A file based distributed storage system comprising:

multiple storage units configured to communicate with each other, each comprising a storage device, a processor, and executable software stored on the storage device, the executable software configured to process file read and write requests on behalf of the distributed storage system;

a switch in communication with the multiple storage units, the switch configured to receive a read request for a file stored on the distributed storage system and to send the read request to one of the multiple storage units wherein each of multiple storage units is operable to monitor in real time a pattern of access to the file, a latency to access each copy of the file, and content included in the file, such that a block cache module will perform K packet read aheads, where K is calculated using a current read rate and a current latency of an access link; and

location metadata necessary to identify the location of a plurality of files stored on the storage system, wherein the location metadata is distributed across a subset of the multiple storage units, each storage unit in the subset storing a portion of the location metadata that is different at least in part from portions stored on other storage units in the subset.

76. (Previously presented) The file based distributed storage system of Claim 75, the location metadata corresponding to a hierarchical directory structure.

77. (Previously presented) The file based distributed storage system of Claim 75, the location metadata including a first metadata structure and a second metadata structure used to locate plurality of files, the first directory metadata structure referencing the second directory metadata structure, and the second directory metadata structure referencing a data location file.

78. (Previously presented) The file based distributed storage system of Claim 75, wherein the file's content data is stored in parity-protected stripes across a set of storage units in the distributed storage system, wherein the location metadata comprises information for locating the striped content data and its associated parity data, and wherein the executable software is configured to use parity data to reconstruct data stored at an inaccessible data location.

79. (Previously presented) The file based distributed storage system of Claim 75, wherein the file's content data is mirrored, wherein the location metadata comprises information for locating mirrored copies of the content data, and wherein the executable software is configured to locate a mirrored copy of data stored at an inaccessible data location.

80. (Previously presented) The file based distributed storage system of Claim 75, wherein a data location file is stored on at least two, but not all, storage units in the distributed storage system, and wherein the executable software is configured to locate an available copy of the data location file.

81. (Previously presented) The file based distributed storage system of Claim 75, wherein a data location file is stored on a storage unit that does not store the file's content data, and wherein the executable software is configured to receive the data location file from the storage unit that does not store the file's content data and to assemble the requested file.

82. (Previously presented) The file based distributed storage system of Claim 75, wherein data location files are mirrored, wherein mirrored copies are stored on different, but not all, storage units in the distributed storage system, and wherein a portion of each directory metadata stores references to the mirrored copies of a data location file.

83. (Previously presented) The file based distributed storage system of Claim 75, wherein the executable software module is configured to assemble the requested file from a data location table comprising, for each data block of the file's content data, a storage unit identifier and a corresponding block address, the storage unit identifier identifying a storage unit in the distributed storage system and the corresponding block address indicating an address within the identified storage unit.

84. (Currently Amended) A distributed file storage system comprising:

a first storage unit comprising:

a first storage device configured to store data;

a first processor in communication with the first storage device configured to execute at least the following software modules stored on the first storage device:

a first data allocation manager module configured to locate, in response to read requests, storage locations in the distributed file storage system corresponding to user files, wherein for each user file there is a separate address file comprising storage locations for file content data, the separate address file stored on at least one, but not all, storage units in the distributed file storage system;

a first data cache module configured to cache data associated with read requests;

a first local data manager module configured to manage data at storage locations local to the first storage unit;

a first remote data manager module configured to communicate with other storage units that store data at storage locations remote to the first storage unit; and

a first storage device module configured to operate the first storage device;

a second storage unit comprising:

a second storage device configured to store data;

a second processor in communication with the second storage device configured to execute at least the following software modules stored on the first storage device:

a second data allocation manager module configured to locate, in response to read requests, storage locations in the distributed file storage system corresponding to user files, wherein for each user file there is a separate address file comprising storage locations for file content data, the separate address file stored on at least one, but not all, storage units in the distributed file storage system;

a second data cache module configured to cache data associated with read requests;

a second local data manager module configured to manage data at

storage locations local to the second storage unit;

a second remote data manager module configured to communicate with other storage units that store data at storage locations remote to the second storage unit; and

a second storage device module configured to operate the second storage device;

a third storage unit comprising:

a third storage device configured to store data;

a third processor in communication with the third storage device configured to execute at least the following software modules stored on the first storage device:

a third data allocation manager module configured to locate, in response to read requests, storage locations in the distributed file storage system corresponding to user files, wherein for each user file there is a separate address file comprising storage locations for file content data, the separate address file stored on at least one, but not all, storage units in the distributed file storage system;

a third data cache module configured to cache data associated with read requests;

a third local data manager module configured to manage data at storage locations local to the third storage unit;

a third remote data manager module configured to communicate with other storage units that store data at storage locations remote to the third storage unit; and

a third storage device module configured to operate the third storage device; wherein each of first storage unit, the second storage unit, and the third storage unit are configured to respond to read requests on behalf of the system; and

a switch in communication with at least the first, second and third storage units, the switch configured to receive a read request for a file stored on the distributed file storage system and to send the read request to at least one of the first, second and third storage units wherein each of storage units is operable to monitor in real time a pattern of access to the file, a latency to access each copy of the file, and content included in the file, such that a block cache module will perform K packet read aheads, where K is calculated using a current read rate and a current latency of an

access link.

85. (Currently Amended) A distributed file storage system comprising:

multiple storage units in communication with each other, each configured to process file read requests on behalf of the distributed file storage system, and each comprising:

a storage device configured to store data;

a processor configured to execute at least the following software modules stored on the storage device:

a data allocation manager module configured to locate, in response to read requests, storage locations in the distributed file storage system corresponding to user files, wherein for each user file there is a separate address file comprising storage locations for file content data, the separate address file stored on at least one, but not all, of the multiple storage units;

a data cache module configured to cache data associated with read requests;

a local data manager module configured to manage data at local storage locations;

a remote data manager module configured to communicate with other storage units that store data at storage locations remote to the storage unit; and

a storage device module configured to operate the storage device; and

a switch in communication with the multiple storage units, the switch configured to receive a read request for a file stored on the distributed file storage system and to send the read request to one of the multiple storage units wherein each of multiple storage units is operable to monitor in real time a pattern of access to the file, a latency to access each copy of the file, and content included in the file, such that a block cache module will perform K packet read aheads, where K is calculated using a current read rate and a current latency of an access link.

86. (Previously presented) The distributed file storage system of Claim 85, wherein the data allocation manager module is further configured to change the allocation of the blocks or the file metadata across the storage modules of the storage system as the storage or retrieval demands

change over time.

87. (Currently Amended) A distributed file storage system comprising:

multiple storage units configured to communicate with each other, each storage unit comprising a storage device, a processor, and at least one executable software module stored on the storage device;

the software module configured to:

write files in a distributed file system;

in response to a request to write a file, initiate the storage of the file's content data into the distributed storage system and to track data location information for locating each portion of the file's content data in a data location file; and

initiate the storage of at least one portion of the file on a different storage unit than other portions of the file; and

wherein the data location information for the file comprises data locations on different storage units; and

a switch in communication with the multiple storage units, the switch configured to receive a read request for a file stored on the distributed file storage system and to send the read request to one of the multiple storage units wherein each of multiple storage units is operable to monitor in real time a pattern of access to the file, a latency to access each copy of the file, and content included in the file, such that a block cache module will perform K packet read aheads, where K is calculated using a current read rate and a current latency of an access link;

88. (Previously presented) The file based distributed storage system of Claim 87, wherein the request to write the file is at least one of the following: a request to write a new file, a request to modify an existing file, and a request to move the file to a new location.

89. (Previously presented) The file based distributed storage system of Claim 87, wherein to track data location information comprises at least one of the following: to store data location information into a new data location file, and to update data location information in an existing

data location file.

90. (Currently Amended) The file based distributed storage system of Claim 87, wherein the file's content data is stored in parity-protected stripes across a set of storage units in the distributed file storage system, and wherein the data location file comprises information for locating the striped content data and its associated parity data.

91. (Previously presented) The file based distributed storage system of Claim 87, wherein the file's content data is mirrored, and wherein the data location file comprises information for locating mirrored copies.

92. (Previously presented) The file based distributed storage system of Claim 87, wherein content data for at least one file is stored with a different kind of protection than content data for another file, and wherein content data is stored with at least one of the following kinds of protection: no protection, mirror protection, parity protection, and array mirror split protection.

93. (Previously presented) The file based distributed storage system of Claim 87, wherein content data for at least one file is stored with a different level of protection than content data for another file, and wherein content data is stored with at least one of the following levels of protection: no protection, mirror protection, multiple mirror protection, parity protection, and multiple parity protection.

94. (Currently Amended) The file based distributed storage system of Claim 87, wherein the data , location file is stored on at least two, but not all, storage units in the distributed file storage system.

95. (Previously presented) The file based distributed storage system of Claim 87, wherein the data location file is stored on a storage unit that does not store the file's content data.

96. (Currently Amended) The file based distributed storage system of Claim 87, wherein data location files are mirrored, and wherein mirrored copies are stored on different, but not all, storage units in the distributed file storage system.

97. (Currently Amended) The file based distributed storage system of Claim 87, wherein the data location file comprises a data location table comprising, for each data block of the file's content data, a storage unit identifier and a corresponding block address, the storage unit identifier identifying a storage unit in the distributed file storage system and the corresponding block address indicating an address within the identified storage unit.

98. (Previously Presented) The distributed file storage system of Claim 1, wherein the second set of file data comprises a next set of file data, wherein the next set of file data is sequential to the first set of file data.